
Feedback-Controlled Sequential Lasso Screening

Yun Wang, Xu Chen and Peter J. Ramadge
 Department of Electrical Engineering
 Princeton University

Abstract

One way to solve lasso problems when the dictionary does not fit into available memory is to first screen the dictionary to remove unneeded features. Prior research has shown that sequential screening methods offer the greatest promise in this endeavor. Most existing work on sequential screening targets the context of tuning parameter selection, where one screens and solves a sequence of N lasso problems with a fixed grid of geometrically spaced regularization parameters. In contrast, we focus on the scenario where a target regularization parameter has already been chosen via cross-validated model selection, and we then need to solve many lasso instances using this fixed value. In this context, we propose and explore a feedback controlled sequential screening scheme. Feedback is used at each iteration to select the next problem to be solved. This allows the sequence of problems to be adapted to the instance presented and the number of intermediate problems to be automatically selected. We demonstrate our feedback scheme using several datasets including a dictionary of approximate size 100,000 by 300,000.

1 Introduction

Sparse dictionary-based representation of data has proved effective in a wide range of applications in computer vision, machine learning, signal processing and statistics [10, 35]. A sparse representation uses a dictionary $D \in \mathbb{R}^{n \times p}$, with columns $\{\mathbf{d}_i\}_{i=1}^p$ called *features*, to represent a data point \mathbf{y} as a linear combination $\sum_{i=1}^p w_i \mathbf{d}_i$ of a relatively small number of the features. Thus for many i , we require $w_i = 0$. One way to obtain such representations is by solving the lasso problem [24]:

$$\min_{\mathbf{w} \in \mathbb{R}^p} \quad \frac{1}{2} \|\mathbf{y} - D\mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_1, \quad (1)$$

where $\lambda > 0$ is a regularization parameter.

In many of today's sparse coding applications, the dictionary can be very large, both because the data dimension is large ([36, 28, 23, 14]) and because an adequate representation requires many features ([36, 22, 1, 34, 14]). For example, in music genre classification the authors of [7] employed scattering representations [1, 2] of short segments of music data, resulting in dictionaries of size $n = 199$ and $p = 12,000$ (first order scattering), and $n = 11,726$ and $p = 12,000$ (second order scattering). Once dictionaries reach this size, solving (1) can become a bottleneck in the overall computation. More importantly, for significantly larger dictionaries, fitting the dictionary into available memory is a concern.

Several approaches are possible for solving (1) with a large dictionary. For example, assuming the dictionary fits into memory, one could use early termination of an iterative solver (e.g., FISTA [3]) to quickly obtain an approximate solution. Alternatively, one could use a finite number of steps of a sequential greedy method (e.g., OMP [26]) to approximate a solution of (1). OMP iteratively selects a feature with the largest correlation with the current residual. Searching for this feature is the most time consuming part of the algorithm but only requires part of the dictionary to be loaded at once. To obtain a solution $\hat{\mathbf{w}}$ that is m -sparse, the algorithm needs to hold at most m features in the memory.

However, since the true sparsity is unknown, in general, OMP and similar algorithms will only yield an approximate solution. Both of these approaches lead (in general) to an approximate solution of (1).

Several recent studies have investigated the following alternative approach. For a given \mathbf{y} and λ , use duality to quickly identify a subset of features in D that are certified to have zero weight in a solution $\hat{\mathbf{w}}$ of (1). By removing these features from the dictionary and solving a smaller lasso problem, we can obtain an exact solution to the original problem. In [9] this approach is called SAFE screening, indicating that screening does not reject any needed feature. Just as important, screening can be executed with only a few features loaded into memory at once. Hence it has the potential to significantly reduce the size of dictionary that is provided to a lasso solver. Recent work on safe screening has further developed this idea [39, 37, 8, 29, 33, 11, 19, 20]. The survey [38] reviews these recent methods. A related but distinct form of lasso screening [25] allows false rejection of features with small probability. We do not explicitly consider such screening tests, but our results are also relevant to this approach.

The core idea of SAFE screening is to bound the solution of the dual problem of (1) within a region \mathcal{R} and compute $\mu(\mathbf{d}_i) = \max_{\boldsymbol{\theta} \in \mathcal{R}} \mathbf{d}_i^T \boldsymbol{\theta}$. The value $\mu(\mathbf{d}_i)$ is then used to decide if feature \mathbf{d}_i can be removed. Since these tests are applied once prior to solving (1), we refer to them as “one-shot” tests. A parameter that is often useful in such tests is the value $\lambda_{\max} = \max_i |\mathbf{d}_i^T \mathbf{y}|$. Roughly, this gives a measure of how well a feature in D matches \mathbf{y} . Empirical studies indicate that current one-shot screening tests perform well for normalized values of λ/λ_{\max} of moderate to large size, but performance quickly declines as λ/λ_{\max} falls below 0.4 [38]. To address this problem, one can employ a more complex form of screening known as sequential screening [9, 25]. In sequential schemes, one screens and solves a sequence of N lasso problems with $\lambda_1 > \dots > \lambda_N$. The key point is that for instance (\mathbf{y}, λ_k) , one utilizes the dual solution $\hat{\boldsymbol{\theta}}_{k-1}$ to obtain a region bound on $\hat{\boldsymbol{\theta}}_k$. This is used to apply a one-shot screening test to reduce the dictionary. And a standard solver is then used to solve the reduced problem to find $\hat{\mathbf{w}}_k$.

Existing state-of-the-art sequential screening algorithms include sequential Dome [32], sequential Strong rule [25], sequential Enhanced DPP rule [29] among others [9]. They are situated in the context of model selection. In this setting, all solutions to the sequence of problems are of interest, and the sequence of regularization parameters is thus fixed a priori (typically taking a log-grid of values). Sequential screening is then conducted along this predetermined sequence to expedite this parameter tuning process.

In contrast, we focus on another distinct application context where the best regularization parameter, denoted by λ_t , has already been chosen via cross-validated model selection, and then we need to solve many lasso instances using this fixed value. Of our particular interest, is λ_t/λ_{\max} around 0.1 since many applications of the sparse representation framework have found this range to be most helpful. For example, [39] found that λ_t/λ_{\max} in the range $[0.1, 0.2]$ maximizes SVM classification accuracy for the handwritten digit recognition problem. In music genre classification, the authors of [7] selected λ_t/λ_{\max} in the range $[0.1, 0.15]$ via cross-validation to maximize classification accuracy using the SRC classifier [36]. Once the proper λ_t to use is known, the problem of efficient sequential screening can be very different from the sequential screening that is targeted for the model selection purpose: the solution at λ_t is the only one that is of interest, and all other problem instances are merely way points. We are thus given the freedom to design the sequence of regularization parameters $\{\lambda_k\}_{k=1}^N$ with $\lambda_N = \lambda_t$ that specifically target a single problem instance at λ_t . In this scenario, the open-loop sequential screening schemes using a geometrically spaced sequence $\{\lambda_k\}_{k=1}^N$, which is determined before any problem in the sequence is solved, may not be a good idea. And it would be an advantage if $\{\lambda_k\}_{k=1}^N$ could be tuned individually to each instance $(D, \mathbf{y}, \lambda_t)$.

This paper makes two contributions. First, we design a feedback mechanism to adaptively select the sequence $\{\lambda_k\}_{k=1}^N$ and N for sequential screening. The feedback mechanism automatically selects the next value λ_k as a function of the results seen in previous steps of the sequential screening process. It also determines when to stop, and hence automatically selects N . We call our feedback controlled sequential screening scheme, data-adaptive sequential screening (DASS). DASS has the advantage that N and $\{\lambda_k\}_{k=1}^N$ are automatically adapted during the screening process to the particular instance $(D, \mathbf{y}, \lambda_t)$. In addition, the feedback mechanism bounds the region diameter used for one-shot screening at each iteration by a value set by the user. Second, we examine the effects of the inevitable errors that accrue in obtaining the solutions (and dual solutions) of the intermediate lasso

problems during sequential screening. At each step, sequential screening assumes exact knowledge of the previous dual solution. However, practical lasso solvers introduce a small error. In the context of classification, we show that the DASS scheme is robust to these errors.

We give required background in §2. Then we introduce DASS and examine its properties in §3. We show that DASS ensures the diameters of the regions used for the one-shot screening are bounded by a user selected value, and the number of intermediate lasso problems is also bounded. We then address the issue of inaccuracies in lasso solutions. §4 discusses the performance of DASS on a selection of datasets, and we conclude in §5. Proofs are given in the supplementary material.

2 Background

The one-shot region tests discussed in the introduction are based on the Lagrangian dual of (1):

$$\begin{aligned} \max_{\boldsymbol{\theta} \in \mathbb{R}^n} \quad & 1/2 \|\mathbf{y}\|_2^2 - \lambda^2/2 \|\boldsymbol{\theta} - \mathbf{y}/\lambda\|_2^2 \\ \text{s.t.} \quad & |\mathbf{d}_i^T \boldsymbol{\theta}| \leq 1 \quad \forall i = 1, 2, \dots, p, \end{aligned} \quad (2)$$

with the solutions $\hat{\mathbf{w}} \in \mathbb{R}^p$ of (1) and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^n$ of (2) related by:

$$\mathbf{y} = D\hat{\mathbf{w}} + \lambda\hat{\boldsymbol{\theta}}, \quad \mathbf{d}_i^T \hat{\boldsymbol{\theta}} = \begin{cases} \text{sign } \hat{w}_i, & \text{if } \hat{w}_i \neq 0; \\ \gamma \in [-1, 1] & \text{if } \hat{w}_i = 0. \end{cases} \quad (3)$$

This dual problem has been extensively discussed in the literature, [38, 39]. The set of feasible points \mathcal{F} of the dual problem is the nonempty, closed, convex polyhedron formed by the intersection of the finite set of closed half spaces $\mathbf{d}^T \boldsymbol{\theta} \leq 1$ for $\mathbf{d} \in \{\pm \mathbf{d}_i\}_{i=1}^p$. The objective in (2) is maximized by the projection $\hat{\boldsymbol{\theta}}$ of \mathbf{y}/λ onto \mathcal{F} . This is the *unique* point satisfying [15, §3.1]:

$$(\mathbf{y}/\lambda - \hat{\boldsymbol{\theta}})^T (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \leq 0, \quad \text{for each } \boldsymbol{\theta} \in \mathcal{F}. \quad (4)$$

Let $\lambda_{\max} = \max_{i=1}^p |\mathbf{d}_i^T \mathbf{y}|$. For $\lambda \geq \lambda_{\max}$, $\hat{\boldsymbol{\theta}}(\lambda) = \mathbf{y}/\lambda \in \mathcal{F}$, with $\hat{\boldsymbol{\theta}}(\lambda_{\max}) = \mathbf{y}/\lambda_{\max}$ on the boundary of \mathcal{F} . For $\lambda < \lambda_{\max}$, \mathbf{y}/λ moves away from \mathcal{F} and $\hat{\boldsymbol{\theta}}(\lambda)$ is its unique projection onto the boundary of \mathcal{F} .

To form a one-shot test, one bounds the dual solution $\hat{\boldsymbol{\theta}}$ in a compact set \mathcal{R} , and computes $\mu(\mathbf{d}_i) = \max_{\boldsymbol{\theta} \in \mathcal{R}} \mathbf{d}_i^T \boldsymbol{\theta}$. By (3), if a feature \mathbf{d}_i satisfies $\mu(\mathbf{d}_i) < 1$ and $\mu(-\mathbf{d}_i) < 1$, then $\hat{w}_i = 0$; see [9, 38]. The bounding region \mathcal{R} can be selected as the intersection of a sphere $S(\mathbf{q}, r) = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \mathbf{q}\|_2 \leq r\}$, with center \mathbf{q} and radius r , and m half spaces, i.e., $\mathcal{R} = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \mathbf{q}\|_2 \leq r, \mathbf{n}_i^T \boldsymbol{\theta} \leq c_i, i = 1, \dots, m\}$. Then $\mu(\mathbf{d}_i) = \max_{\boldsymbol{\theta} \in \mathcal{R}} \mathbf{d}_i^T \boldsymbol{\theta}$ is obtained by solving a convex program, with closed form solutions available for $m = 0, 1, 2$ [38, 31].

Sequential screening screens and solves a sequence of problems with $\lambda_1 > \dots > \lambda_N$ until $\lambda_N = \lambda_t$. In an open-loop scheme, N and $\{\lambda_k\}_{k=1}^N$ are selected before any solution is obtained. For example, fix N , set $\lambda_1 = 0.95\lambda_{\max}$, and space the λ_k geometrically: $\lambda_k = \alpha\lambda_{k-1}$ with $\alpha = (\lambda_t/\lambda_1)^{1/(N-1)}$.

Note that $\hat{\boldsymbol{\theta}}_1$ must lie in the sphere $S(\mathbf{q}_1, r_1)$ with center $\mathbf{q}_1 = \mathbf{y}/\lambda_1$ and radius $r_1 = \|\mathbf{y}/\lambda_1 - \mathbf{y}/\lambda_{\max}\|_2$ [38]. In addition, if \mathbf{d}_{\max} is a feature satisfying $|\mathbf{d}_{\max}^T \mathbf{y}| = \lambda_{\max}$, then $\hat{\boldsymbol{\theta}}_1$ lies in the half space $H_{\max} = \{\boldsymbol{\theta} : \mathbf{d}_{\max}^T \boldsymbol{\theta} \leq 1\}$. So to screen the dictionary at $k = 1$, we use the one-shot screening test for

$$\mathcal{R}_1^0 = S(\mathbf{q}_1, r_1) \cap H_{\max}. \quad (5)$$




Figure 1: Sequential screening. At step k , $\hat{\boldsymbol{\theta}}_{k-1}$ defines a spherical bound (solid red curve) on $\hat{\boldsymbol{\theta}}_k$ (eq. (6)). The (blue dash-dot) hyperplane $\mathbf{n}_{k-1}^T \boldsymbol{\theta} \leq c_{k-1}$ (eq. (7)) separates \mathcal{F} from \mathbf{y}/λ_{k-1} . Hence $\hat{\boldsymbol{\theta}}_k$ is contained in the dome \mathcal{R}_k^0 (eq. (8)) formed within the sphere and above the hyperplane. The solid line (black) is the hyperplane of an active feature at $\hat{\boldsymbol{\theta}}_{k-1}$. The large hemisphere (green) is the DPP bound $\|\hat{\boldsymbol{\theta}}_k - \hat{\boldsymbol{\theta}}_{k-1}\|_2 \leq 1/\lambda_k - 1/\lambda_{k-1}$ used in [29]. The diagram is drawn to scale.

Then we solve the reduced lasso problem to obtain $\hat{\boldsymbol{\theta}}_1$ (see (3)).

At step k , we make use of the dual solution $\hat{\boldsymbol{\theta}}_{k-1}$ for instance $(\mathbf{y}, \lambda_{k-1})$. Since $\hat{\boldsymbol{\theta}}_{k-1}$ is dual feasible, $\hat{\boldsymbol{\theta}}_k$ is contained in the sphere $S(\mathbf{q}_k, r_k)$ with

$$\mathbf{q}_k = \mathbf{y}/\lambda_k, \quad r_k = \|\mathbf{y}/\lambda_k - \hat{\boldsymbol{\theta}}_{k-1}\|_2. \quad (6)$$

In addition, since $\hat{\boldsymbol{\theta}}_{k-1}$ is the projection of \mathbf{y}/λ_{k-1} onto \mathcal{F} , by (4) we have $\forall \boldsymbol{\theta} \in \mathcal{F}, (\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1})^T \boldsymbol{\theta} \leq (\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1})^T \hat{\boldsymbol{\theta}}_{k-1}$. Since $0 \in \mathcal{F}$, the right hand side of this inequality is nonnegative. Hence \mathcal{F} lies in the closed half space $H_k = \{\boldsymbol{\theta} : \mathbf{n}_{k-1}^T \boldsymbol{\theta} \leq c_{k-1}\}$ with

$$\mathbf{n}_{k-1} = \frac{\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1}}{\|\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1}\|_2}, \quad c_{k-1} = \mathbf{n}_{k-1}^T \hat{\boldsymbol{\theta}}_{k-1}. \quad (7)$$

In particular, $\hat{\boldsymbol{\theta}}_k \in \mathcal{F}$, so $\hat{\boldsymbol{\theta}}_k$ is contained in the region

$$\mathcal{R}_k^0 = S(\mathbf{q}_k, r_k) \cap H_k. \quad (8)$$

This is illustrated in Fig. 1. We now use the one-shot screening test for \mathcal{R}_k^0 to screen the dictionary. Then solve the reduced lasso problem to obtain $\hat{\boldsymbol{\theta}}_k$.

3 Feedback Controlled Sequential Screening

Under the assumption of exact computation, we now develop our sequential screening feedback rule and analyze its properties. The issue of approximate lasso solution is examined later. The diameter of a set \mathcal{R} , denoted by $\text{diam}(\mathcal{R})$, is the maximum distance between any two points in \mathcal{R} . The following preliminary lemma will be useful.

Lemma 1. *If $\mathcal{R} = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \mathbf{q}\|_2 \leq r, \mathbf{n}^T \boldsymbol{\theta} \leq c\}$ has a nonempty interior, then*

$$\text{diam}(\mathcal{R}) = \begin{cases} 2\sqrt{r^2 - (\mathbf{n}^T \mathbf{q} - c)^2}, & \text{if } \mathbf{q} \notin \mathcal{R}; \\ 2r, & \text{otherwise.} \end{cases}$$

We next show that if $S(\mathbf{q}_k, r_k)$, H_k and \mathcal{R}_k^0 are constructed via (6), (7), and (8), then $\mathbf{q}_k \notin \mathcal{R}_k^0$.

Lemma 2. *For \mathcal{R}_k^0 constructed via (6), (7), and (8), $\mathbf{q}_k \notin \mathcal{R}_k^0$.*

Combining Lemmas 1 and 2 yields the following result.

Proposition 1. *For \mathcal{R}_k^0 constructed via (6), (7), and (8),*

$$\text{diam}(\mathcal{R}_k^0) = 2 \left(\frac{1}{\lambda_k} - \frac{1}{\lambda_{k-1}} \right) (\mathbf{y}^T (I - \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T) \mathbf{y})^{1/2}. \quad (9)$$

The region \mathcal{R}_k^0 provides an initial bound for the one-shot test applied at step k . More half space constraints can be used to form a final bounding region \mathcal{R}_k . For example, the one-shot test THT in [33, 38, 30] does this by adding a second appropriately selected half space constraint from the dual problem. No matter how, or how many, additional half space constraints are added, (9) provides an upper bound on the diameter of the final bounding region. Thus we have the following corollary.

Corollary 1. *If \mathcal{R}_k is obtained by adding a finite number of additional half space constraints to \mathcal{R}_k^0 , then $\text{diam}(\mathcal{R}_k)$ is bounded above by the expression in (9).*

Proposition 1 suggests a rule for selecting λ_k so that $\text{diam}(\mathcal{R}_k) \leq R$ where $R > 0$ is a user specified parameter. Using (9), this is achieved by selecting λ_k so that:

$$\frac{1}{\lambda_k} = \frac{1}{\lambda_{k-1}} + \frac{1/2R}{\sqrt{\mathbf{y}^T (I - \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T) \mathbf{y}}} \quad (10)$$

Algorithm 1: DASS

Input: A lasso instance $(D, \mathbf{y}, \lambda_t)$, and scalar $R > 0$

Output: Solution $\hat{\mathbf{w}}_t$ of the lasso instance

```

1 Initialize  $N = 0, k = 1$ , and  $\lambda_1$  with
   $\lambda_t < \lambda_1 < \lambda_{\max}$ ;
2 if  $k=l$  then
3   Use region (5) to one-shot screen  $D$ 
   for lasso instance  $(\mathbf{y}, \lambda_1)$ ;
4   Solve the reduced lasso for  $\hat{\mathbf{w}}_1$  and  $\hat{\boldsymbol{\theta}}_1$ ;
5 while  $\lambda_k > \lambda_t$  do
6   Set  $k=k+1$  and use (10) to obtain  $\lambda_k$ ;
7   if  $\lambda_k \leq \lambda_t$  then
8     Set  $N = k$  and  $\lambda_k = \lambda_t$ ;
9   Use  $\hat{\boldsymbol{\theta}}_{k-1}$  and region (8) to one-shot
   screen  $D$  for lasso instance  $(\mathbf{y}, \lambda_k)$ ;
10  Solve the reduced lasso for  $\hat{\mathbf{w}}_k$  and  $\hat{\boldsymbol{\theta}}_k$ ;
11 return  $\hat{\mathbf{w}}_N$ ;
```

sequence $\{\lambda_k\}_{k=1}^N$ be selected using the DASS algorithm. The DASS sequence is guaranteed to terminate after a finite number of steps N , and $\text{diam}(\mathcal{R}_k) \leq R, k = 2:N$. In addition, if the dual regularization path $\hat{\boldsymbol{\theta}}(\lambda)$ is bounded, i.e., there exists $C > 0$, such that $\|\hat{\boldsymbol{\theta}}(\lambda)\| \leq C$ for all $\lambda > 0$, then

$$N \leq 1 + \frac{\log(1/\lambda_t)}{\log(1 + R/2C)}.$$

The assumption in Theorem 1 that $\hat{\boldsymbol{\theta}}(\lambda)$ is bounded holds if \mathbf{y} is in the range of D (supplementary material). It always holds if $\text{rank}(D) = n$ [40]. We show the number of iterations required for the algorithm to terminate in real problem instances in §4.

In practice, it is more realistic to assume we have an approximation $\tilde{\boldsymbol{\theta}}_{k-1}$ to $\hat{\boldsymbol{\theta}}_{k-1}$, with $\|\hat{\boldsymbol{\theta}}_{k-1} - \tilde{\boldsymbol{\theta}}_{k-1}\|_2 \leq \rho$. For simplicity we assume ρ does not depend on k , but this is not essential. We first verify that the DASS algorithm still terminates after a finite number of steps and give a bound on the number of steps required. We state this as a corollary to Theorem 1.

Corollary 2. Suppose we have an approximate dual solution $\tilde{\boldsymbol{\theta}}_k$, with $\|\hat{\boldsymbol{\theta}}_k - \tilde{\boldsymbol{\theta}}_k\|_2 \leq \rho, k = 1:N$. Under this assumption, the DASS sequence terminates with $N \leq 1 + \frac{\log(1/\lambda_t)}{\log(1 + R/2(C+\rho))}$.

In many classification applications, the lasso regression is often used as a feature extraction step. For each data point \mathbf{y} to be classified, lasso regression is first used to compute the solution $\hat{\mathbf{w}}$. Then as a feature vector, $\hat{\mathbf{w}}$ is fed into a classifier, e.g., SVM, SRC. We analyze in this classification setting, the effect of accumulated inaccuracy in our sequential screening process. Consider using an approximate solution $\tilde{\mathbf{w}} = \hat{\mathbf{w}} + \mathbf{n}$ for classification. When used in sequential screening, $\tilde{\mathbf{w}}_{k-1}$ gives an approximate dual solution $\tilde{\boldsymbol{\theta}}_{k-1}$, which is used for one-shot screening at step k . So the one-shot test may both fail to reject, and falsely reject features. The error in $\tilde{\mathbf{w}}_k$ thus has two sources: from the lasso solver itself, and from false feature rejection. The latter errors can propagate, leading DASS to fail to give adequate rejection, or to make too many false rejections causing classification performance to suffer. We investigate this in §4.

4 Empirical Tests

We compare DASS with state-of-the-art open-loop sequential screening schemes with a fixed grid of geometrically spaced regularization parameters. These open-loop sequential screening schemes are sequential Dome [32], sequential Strong rule [25] and sequential Enhanced DPP rule [29]. Sequential Dome is the sequential version of the dome test proposed in [37]. The authors of [29]

Assuming \mathbf{y} is not aligned with $\hat{\boldsymbol{\theta}}_{k-1}$ (this holds generically), $\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y} > 0$. Incorporating (10) into a sequential screening scheme yields the DASS algorithm in Algorithm 1. We can initialize $\lambda_1 = 0.95\lambda_{\max}$. At each step, we first use one-shot screening based on the region \mathcal{R}_k^0 , then solve the resulting reduced lasso problem (for $k > 1$, $\hat{\mathbf{w}}_{k-1}$ can be used as a warm start). Hence each problem is solved efficiently.

If we write $s_k = 1/\lambda_k$, then (10) becomes $s_k - s_{k-1} = f(\mathbf{n}_{k-1}^T \mathbf{y})$. Thus s_k is a first order, autoregressive process driven by a nonnegative, nonlinear function of $\mathbf{n}_{k-1}^T \mathbf{y}$. Since \mathbf{n}_{k-1} depends on the solution of the previous instance, this term is providing feedback in the update of s_k . The feedback rule ensures that λ_k decreases at each step; this results in $\lambda_k < \lambda_t$, and termination of the algorithm. We show in Theorem 1 that the algorithm always terminates, give an upper bound on the number of iterations required, and verify that for $k > 1$ the one-shot tests use a region of diameter at most R .

Theorem 1. For $0 < \lambda_t < \lambda_{\max}$, let N and the




Figure 2: $\lambda_t/\lambda_{\max} \approx 0.1$ is of particular interest as cross validation shows this range is most helpful for many classification problems. (From left to right): handwritten digit recognition on MNIST [16], music genre classification on GTZAN [27], object recognition on COIL [21] and text categorization on RCV1 [6].

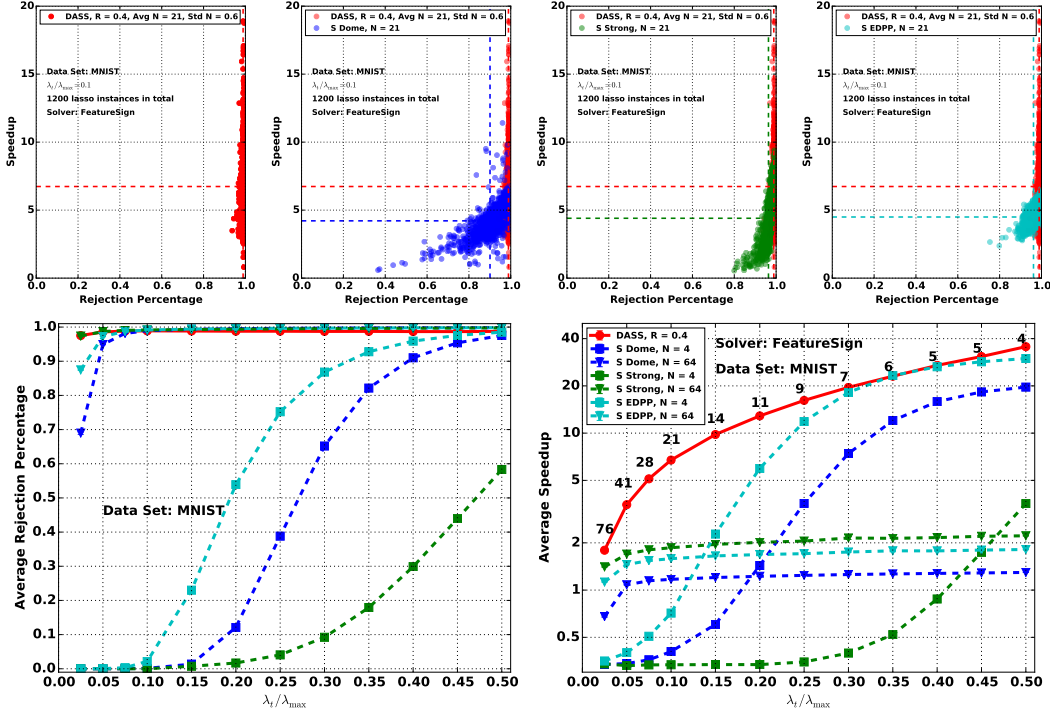


Figure 3: **Comparison of DASS with sequential Dome, sequential Strong rule and sequential Enhanced DPP rule on MNIST.** (Top): Scatter plot of speedup versus rejection percentage for lasso instances (D, y) with $\lambda_t/\lambda_{\max} = 0.1$. (Bottom): Average rejection and speedup for $\lambda_t/\lambda_{\max} \in [0.025, 0.5]$. The average N for DASS is shown on the speedup curve.

commented that “for the DOME test, it is unclear whether its sequential version exists” and did not include a comparison of their sequential Enhanced DPP rule with sequential Dome. Their claim is not true and we hereby include both algorithms for comparison. Among these four screening algorithms, only the sequential Strong rule in [25] may falsely discard features.

Performance Metrics

1) Rejection Percentage: the ratio of the number of features rejected by screening rules to the number of *all* features.

2) Speedup: the time to solve (1) at λ_t without screening (*one problem*) divided by the total time to screen *plus* solve the reduced lasso problem along the sequence (*N problems*).

Note the difference in our metrics definition compared with [29]. In [29] rejection percentage was defined as the ratio of the number of features discarded by screening rules to the actual number of *zero* features in the ground truth; and speedup was defined as the time to solve the sequence of lasso problems without screening (*N problem*) divided by the total time to screen plus solve the reduced lasso problem along the same sequence (*N problem*). The difference in speedup definition mainly originates from our different motivation and application, as we discussed earlier in §1.

Datasets

(a) MNIST: 28×28 hand-written digit images (training set: 60,000, testing set: 10,000) [16]. Each image is rearranged into a vector and scaled to unit norm. We sample 5,000 training images (500 per digit) to form the dictionary. Test images \mathbf{y} are randomly selected from the testing set.

(b) GTZAN: 100 music clips (30 sec, sampled at 22,050 Hz) for each of ten genres of music [27]. Each clip is divided into 3-sec texture windows (TW) with 50% overlap of adjacent TW. Each TW is then represented using a first order scattering vector of length 199 [1]. We randomly select a dictionary of 12,000 scattering vectors and randomly select test vectors \mathbf{y} from the remaining 8,000.

(c) NYT: 300,000 New York Times articles represented as vectors with respect to a vocabulary of 102,660 words [12]. The i -th entry in vector j gives the occurrence count of word i in document j . After removing documents with very low word counts, 299,752 documents remain. 299,652 of these documents are randomly selected to form the dictionary and from the remaining 100 documents, 65 are selected as test vectors with $0.3 < \lambda_{\max} < 0.9$.

For experiments on MNIST and GTZAN we randomly select 20 dictionaries and for each dictionary we use 60 randomly selected test vectors. Reported results are averaged across selected dictionaries and test vectors. λ_t is set using the scaling invariant ratio $\lambda_t/\lambda_{\max} \in [0, 1]$ with higher value yielding sparser solution. Speedup results using the Feature-sign lasso solver [17] are shown but our experiment indicate consistent results for a variety of other solvers. For example, see the supplementary material for additional results using the FISTA [4] solver.

In an open-loop sequential screening scheme, once λ_t is given and N is selected then the sequence $\{\lambda_k\}_{k=1}^N$ is fixed. We claim a key attribute of DASS is that it adapts $\{\lambda_k\}_{k=1}^N$ to each individual problem: $(D, \mathbf{y}, \lambda_t)$. We test the effectiveness of this adaptation by solving 1200 problem instances with $\lambda_t/\lambda_{\max} = 0.1$ for both MNIST and GTZAN. We then plot the scatter plot of speedup versus rejection percentage for these 1200 problem instances. See the top rows of Fig. 3 and Fig. 4. For MNIST, the average N used by DASS ($R = 0.4$) is 21 with standard error of 0.6. DASS successfully pushes both the rejection percentage and speedup distribution to higher ranges compared with open-loop schemes using $N = 21$.

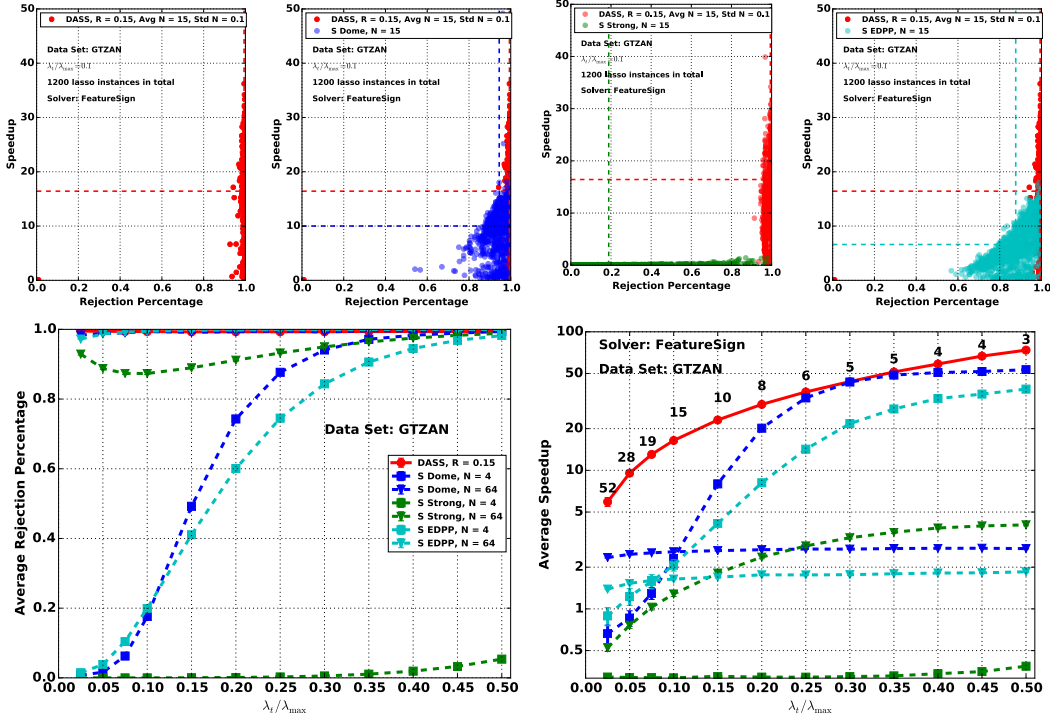


Figure 4: Comparison of DASS with sequential Dome, sequential Strong rule and sequential Enhanced DPP rule on GTZAN. (Top): Scatter plot of speedup versus rejection percentage for lasso instances (D, \mathbf{y}) with $\lambda_t/\lambda_{\max} = 0.1$. (Bottom): Average rejection and speedup for $\lambda_t/\lambda_{\max} \in [0.025, 0.5]$. The average N for DASS is shown on the speedup curve.

DASS also exhibits robust average rejection percentage and speedup as λ_t/λ_{\max} varies (bottom rows Fig. 3 and 4). For MNIST, DASS consistently beats or rivals the open-loop schemes as N ranges from 4 to 64. Open-loop schemes with $N = 64$ are close to DASS in rejection, but its average speedup is only around 2 or less for most values of λ_t/λ_{\max} . The DASS speedup curve gives a nice outer envelope across the set of open-loop speedup curves. The average N used by DASS corresponding to each value of λ_t/λ_{\max} is given on top of the DASS speedup curve (red). As λ_t/λ_{\max} decreases, DASS increases N on a case-by-case basis, and this is done without user intervention. Using a properly cross-validated R , DASS is able to automatically design an efficient sequence $\{\lambda_k\}_{k=1}^N$ for each individual $(D, \mathbf{y}, \lambda_t)$. Similar results can be seen for GTZAN. Since overhead accumulates as N increases, DASS, and open-loop sequential screening schemes, both exhibit less speedup as λ_t/λ_{\max} decreases. Nevertheless, DASS pushes efficient and robust screening down to $\lambda_t/\lambda_{\max} = 0.025$, which includes a range useful in many classification applications.

The DASS scheme is easy to implement, and its one parameter can be selected by cross-validation using one reasonable value of λ_t/λ_{\max} . The scheme then gives strong screening performance across a wide range of values of λ_t/λ_{\max} . A smaller value of the user-specified parameter R , yields a tighter bound on $\text{diam}(\mathcal{R}_k^0)$, and hence stronger rejection, but also a larger value of N . Hence speedup will hit a sweet spot as R decreases. Empirical investigation on the effect of R is given in Fig. 5.




Figure 5: **Effects of the choice of R on DASS performance.** (Left): Average rejection percentage. (Middle): Average speedup. (Right): The average value of N over the 1200 trials.

To examine the effect of “noisy” lasso solutions, we add noise $\mathbf{n}_k \sim N(0, \alpha^2 I)$ to our computed $\hat{\mathbf{w}}_k$, $k = 1:N$, for various values of noise-to-signal ratio (the ratio of the power of \mathbf{n}_k to that of $\hat{\mathbf{w}}_k$). We then hard threshold $\hat{\mathbf{w}}_k + \mathbf{n}_k$ to yield a “corrupted” sparse solution. The impact on rejection fraction, speedup, and classification accuracy is shown for MNIST in Fig. 6. Rejection remains around 0.9 until the noise-to-signal ratio goes above 10^{-3} . Speedup drops from around $5X$ to $1X$ at 10^{-6} . So DASS remains effective for noise-to-signal ratio below 10^{-6} . Within this bound, there is a small penalty for the added corruption, but DASS is robust to the errors. In addition, when $\tilde{\mathbf{w}}_N$ is used for Sparse Representation Classification [36], classification accuracy is also robust to the error in $\tilde{\mathbf{w}}_N$.

Finally, we use the NYT dataset to explore problems with many high dimensional features. From the test vector pool of 100 documents, we select 65 documents, subject to $0.3 < \lambda_{\max} < 0.9$. λ_t/λ_{\max} is set at 0.1. Since the dictionary cannot fit into our computer memory, we run screening by loading small segments of the dictionary into memory at a time. We test DASS with $R = 0.3$. The average resulting values of N from 65 problem instances is 45 with standard error being 1.8. For comparison, we run sequential Dome, sequential Strong rule and sequential Enhanced DPP rule with $N = 65$.

We allocate the same amount of system resources to each algorithm and each problem instance. Sequential Dome and sequential EDPP are not able to complete all 65 problem instances because for some instances they do not reject enough features in the process and run into out of memory




Figure 6: Effects of adding noise to \hat{w}_k in DASS.




Figure 7: Comparison of DASS with sequential Dome, sequential Strong rule and sequential Enhanced DPP rule on NYT. Both DASS and three open-loop sequential screening schemes are applied to 65 lasso instances with $0.3 < \lambda_{\max} < 0.9$ and $\lambda_t/\lambda_{\max} = 0.1$. The colored + sign marks the average performance for each algorithm.

error. We thus include the completion rate metric in Fig. 7 defined as the percentage of problem instances that are completed given the same system resources. DASS and sequential Strong rule take the lead in this metric and are able to complete all problem instances. For problem instances that are completed by each algorithm, we provide a scatter plot of total running time versus number of remaining features in Fig. 7. No speedup information is shown since we cannot solve these problems without screening. The results show that DASS is on a par with sequential Strong rule and they significantly outperform sequential Dome and sequential EDPP rules: sequential Strong rule is able to reject more but DASS is able to finish in less runtime. Note however that the Strong rule may commit false rejections.

5 Conclusion

We have shown analytically and empirically that feedback controlled sequential screening yields improved performance. Given a cross-validated target λ_t , the DASS rule tailors an individualized sequence $\{\lambda_k\}_{k=1}^N$ on the fly which is optimized for each particular problem instance (D, y, λ_t) . This has demonstrated significantly greater robustness with respect to the inherent variability in the lasso problem.

Our feedback idea extends to the DPP bound used in [29]. The selection rule $\lambda_k = (1/2R + \lambda_{k-1}^{-1})^{-1}$ yields a feedback controlled version of sequential DPP [29]. This uses a looser bound than \mathcal{R}_k^0 (see Fig. 1). Our results also suggest that it may be possible to use feedback to control sequential screening in more targeted ways. For example, by varying the parameter R in response to previously solved instances it may be possible to regulate the size of the dictionary after screening. This would be useful in situations of a large number of features and limited memory.

References

- [1] J. Andén and S. Mallat. Multiscale scattering for audio classification. *Proc. of the ISMIR Conference*, 2011.
- [2] J. Andén and S. Mallat. Deep scattering spectrum. *IEEE Trans. on Signal Processing*, 2013.
- [3] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2(1):183–202, 2009.
- [4] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. IMAGING SCIENCES*, 2(1):183–202, 2009.
- [5] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge Univ Pr, 2004.
- [6] D. Cai and X. He. Manifold adaptive experimental design for text categorization. *IEEE Transactions on Knowledge and Data Engineering*, April 2011.
- [7] Xu Chen and P. J. Ramadge. Music genre classification using multiscale scattering and sparse representations. In *Proc. Annual Conference on Information Sciences and Systems*, 2013.
- [8] L. Dai and K. Pelckmans. An ellipsoidal based, two-stage screening test for bpdn. In *20th European Signal Processing Conference*, Aug. 2012.
- [9] L. El Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination in sparse supervised learning. *Pacific Journal of Optimization*, 8(4):667–698, 2012.
- [10] M. Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, 2010.
- [11] O. Fercoq, A. Gramfort, and J. Salmon. Mind the duality gap: safer rules for the lasso. In *International Conference on Machine Learning*, 2015.
- [12] A. Frank and A. Asuncion. UCI Machine Learning Repository, University of California, Irvine, School of Information and Computer Sciences, 2010.
- [13] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(6):643–660, 2002.
- [14] Y. He, K. Kavukcuoglu, Y. Wang, A. Szlam, and Y. Qi. Unsupervised feature learning by deep sparse coding. In *SIAM International Conference on Data Mining*, pages 902–910. SIAM, 2014.
- [15] J.-B. Hiriart-Urruty and C. Lemarechal. *Fundamentals of Convex Analysis*. Springer, 2001.
- [16] Y. LeCun and C. Cortes. The MNIST database of handwritten digits, 1998.
- [17] H. Lee, A. Battle, R. Raina, and A.Y. Ng. Efficient sparse coding algorithms. In *Advances in Neural Information Processing Systems*, volume 19, page 801, 2007.
- [18] K. C. Lee, J. Ho, and D. J. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, pages 684–698, 2005.
- [19] E. Ndiaye, O. Fercoq, A. Gramfort, and J. Salmon. Gap safe screening rules for sparse multi-task and multi-class models. In *Advances in Neural Information Processing Systems*, 2015.
- [20] E. Ndiaye, O. Fercoq, A. Gramfort, and J. Salmon. Gap safe screening rules for sparse-group lasso. In *Advances in Neural Information Processing Systems*, 2016.
- [21] S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (coil-100). *Techn. Rep. No. CUCS-006-96, dept. Comp. Science, Columbia University*, 1996.
- [22] Y. Panagakis, C. Kotropoulos, and G.R. Arce. Music genre classification using locality preserving non-negative tensor factorization and sparse representations. *Proc. of the ISMIR Conference*, 2009.
- [23] S. Prasad, P. Melville, A. Banerjee, and V. Sindhwani. Emerging topic detection using dictionary learning. In *ACM Conference on Information and Knowledge Management*, 2011.
- [24] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc B.*, 58(1):267–288, 1996.
- [25] R. Tibshirani, J. Bien, J. Friedman, T. Hastie, N. Simon, J. Taylor, and R. J. Tibshirani. Strong rules for discarding predictors in lasso-type problems. *Arxiv preprint arXiv:1011.2234*, 2010.
- [26] J.A. Tropp and A.C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *Information Theory, IEEE Transactions on*, 53(12):4655–4666, Dec 2007.
- [27] G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *IEEE Trans. Speech and Audio Processing*, 10(5), 2002.
- [28] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, H. Mobahi, and Y. Ma. Towards a practical face recognition system: robust alignment and illumination by sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011.
- [29] J. Wang, P. Wonka, and J. Ye. Lasso screening rules via dual polytope projection. *Journal of Machine*

- Learning Research*, to appear.
- [30] Y. Wang. *Feature Screening for the Lasso*. PhD thesis, Princeton University, Sept. 2015.
 - [31] Y. Wang and P. J. Ramadge. The symmetry of a simple optimization problem in lasso screening. *arXiv:1608.06014 [cs.LG]*, 2016.
 - [32] Y. Wang, Z. J. Xiang, and P. J. Ramadge. Lasso screening with a small regularization parameter. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Jun. 2013.
 - [33] Y. Wang, Z.J. Xiang, and P.J. Ramadge. Tradeoffs in improved screening of lasso problems. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Jun. 2013.
 - [34] Yun Wang, Xu Chen, and P. J. Ramadge. Sparse representation classification via sequential lasso screening. In *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*, pages 1001–1004, Dec. 2013.
 - [35] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. Huang, and S. Yan. Sparse representation for computer vision and pattern recognition. *Proceedings of the IEEE*, 98(6):1031–1044, 2010.
 - [36] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 31(2):210–227, 2009.
 - [37] Z. J. Xiang and P. J. Ramadge. Fast lasso screening tests based on correlations. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, 2012.
 - [38] Z. J. Xiang, Yun Wang, and P. J. Ramadge. Screening tests for lasso problems. *IEEE Trans. Pattern Analysis and Machine Intelligence*, DOI: 10.1109/TPAMI.2016.2568185., 12 May 2016.
 - [39] Z. J. Xiang, H. Xu, and P. J. Ramadge. Learning sparse representations of high dimensional data on large scale dictionaries. In *Advances in Neural Information Processing Systems*, 2011.
 - [40] Pingmei Xu and P. J. Ramadge. Three structural results on the lasso. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, 2013.

Supplementary Material

Proof of Lemma 1

Proof. We want to find the solution to $\min_{\theta_1, \theta_2} -1/2 \|\theta_1 - \theta_2\|_2^2$ subject to: $\|\theta_k - \mathbf{q}\|_2^2 \leq r^2$ and $\mathbf{n}^T \theta_k \leq c$, $k = 1, 2$. Form the Lagrangian

$$L = -1/2 \|\theta_1 - \theta_2\|_2^2 + \sum_{k=1}^2 1/2 \lambda_k (\|\theta_k - \mathbf{q}\|_2^2 - r^2) + \sum_{k=1}^2 \mu_k (\mathbf{n}^T \theta_k - c)$$

The objective function and inequality constraint functions are all convex C^1 functions, and by the assumption of Lemma 1 there exist strictly feasible points. Hence the KKT conditions are necessary and sufficient for optimality [5, P. 244]. Thus for θ_1, θ_2 to be a solution, it is sufficient that these points satisfy the four constraints, the complementary slackness conditions $\lambda_k (\|\theta_k - \mathbf{q}\|_2^2 - r^2) = 0$ and $\mu_k (\mathbf{n}^T \theta_k - c) = 0$, $k = 1, 2$, and that $(dL/d\theta_1)(\theta_1) = (dL/d\theta_2)(\theta_2) = 0$.

We now construct a candidate solution. Suppose \mathbf{q} is not contained in R . Let \mathbf{z} denote the orthogonal projection of \mathbf{q} onto the hyperplane $\mathbf{n}^T \theta = c$. Since $\mathbf{q} \notin R$, $\mathbf{n}^T \mathbf{q} > c$ and $\mathbf{z} = \mathbf{q} - \psi r \mathbf{n}$ for some $0 < \psi \leq 1$. Since \mathbf{z} lies on the plane, we must have $c = \mathbf{n}^T \mathbf{z} = \mathbf{n}^T \mathbf{q} - \psi r$. Hence

$$\psi = (\mathbf{n}^T \mathbf{q} - c)/r.$$

Pick any unit norm vector \mathbf{u} with $\mathbf{n}^T \mathbf{u} = 0$. Then set

$$\theta_1 = \mathbf{z} + r\sqrt{1 - \psi^2} \mathbf{u} \quad \theta_2 = \mathbf{z} - r\sqrt{1 - \psi^2} \mathbf{u}. \quad (11)$$

For these points $\|\theta_k - \mathbf{q}\|_2^2 = \|-r\psi \mathbf{n} \pm r\sqrt{1 - \psi^2} \mathbf{u}\|_2^2 = r^2 \psi^2 + r^2(1 - \psi^2) = r^2$ and $\mathbf{n}^T \theta_k = \mathbf{n}^T (\mathbf{z} \pm \psi r \mathbf{n}) = c$. So both points satisfy the inequality constraints with equality and hence also satisfy the complementary slackness conditions. The final two KKT conditions require

$$D_{\theta_1} L(\theta_1) = -(\theta_1 - \theta_2) + \lambda_1(\theta_1 - \mathbf{q}) + \mu_1 \mathbf{n} = 0 \quad (12)$$

$$D_{\theta_2} L(\theta_2) = -(\theta_2 - \theta_1) + \lambda_2(\theta_2 - \mathbf{q}) + \mu_2 \mathbf{n} = 0 \quad (13)$$

Consider equation (12). Substituting θ_1, θ_2 from (11) yields

$$r\sqrt{1 - \psi^2}(\lambda_1 - 2)\mathbf{u} + (\mu_1 - \lambda_1 \psi r)\mathbf{n} = 0.$$

This can be satisfied by selecting $\lambda_1 = 2$ and $\mu_1 = 2\psi r$. A similar argument for (13) shows that it is satisfied by the selection $\lambda_2 = 2$ and $\mu_2 = 2\psi r$. Thus the candidate solution is indeed a solution. So if $\mathbf{q} \notin \mathcal{R}$, we have $\text{diam}(\mathcal{R}) = \|2r\sqrt{1 - \psi^2} \mathbf{u}\|_2 = 2\sqrt{r^2 - (\mathbf{n}^T \mathbf{q} - c)^2}$.

If $\mathbf{q} \in \mathcal{R}$, then for the candidate solution we take $\theta_1 = \mathbf{q} + r\mathbf{u}$ and $\theta_2 = \mathbf{q} - r\mathbf{u}$. In this case, the candidate points satisfy the spherical bound with equality. The half space constraint is also satisfied: $\mathbf{n}^T \theta_k = \mathbf{n}^T \mathbf{q} \pm r\mathbf{n}^T \mathbf{u} \leq c$, but generally only with inequality. So the candidate solution satisfies all of the constraints and the first two complementary slackness conditions. In this case, substitution of the expressions for θ_1, θ_2 into (12) yields

$$r(\lambda_1 - 2)\mathbf{u} + \mu_1 \mathbf{n} = 0$$

This can be satisfied by selecting $\lambda_1 = 2$ and $\mu_1 = 0$. This also satisfies the corresponding complementary slackness condition. Similarly for (13). So when $\mathbf{q} \in \mathcal{R}$, $\text{diam}(\mathcal{R}) = \|2r\mathbf{u}\|_2 = 2r$. \square

Proof of Lemma 2

Proof. $\mathcal{R}_k^0 = S(\mathbf{q}_k, r_k) \cap H_k$. 1) For $k = 1$, $\mathbf{q}_1 = \mathbf{y}/\lambda_1$, $r_1 = \|\mathbf{y}/\lambda_1 - \mathbf{y}/\lambda_{\max}\|_2$, $H_1 = \{\theta : \mathbf{d}_*^T \theta \leq 1\}$. Since $\lambda_1 < \lambda_{\max}$, $\mathbf{q}_1 = \mathbf{y}/\lambda_1$ is outside the dual feasible set \mathcal{F} and outside the half space H_1 . Thus \mathbf{q}_1 is outside \mathcal{R}_1^0 . 2) For $k = 2, \dots, N$, by construction $\mathcal{R}_k^0 \subset H_k = \{\theta : \mathbf{n}_{k-1}^T \theta \leq c_{k-1}\}$, with \mathbf{n}_{k-1} and c_{k-1} given by equation (7). $\mathbf{q}_{k-1} = \mathbf{y}/\lambda_{k-1} \notin \mathcal{F}$, since $\lambda_{k-1} < \lambda_1 < \lambda_{\max}$. Since $\{\theta : \mathbf{n}_{k-1}^T \theta = c_{k-1}\}$ is a separating hyperplane of \mathcal{R}_{k-1} from \mathcal{F} , $\mathbf{n}_{k-1}^T \mathbf{q}_{k-1} > c_{k-1}$. Then $\mathbf{n}_{k-1}^T \mathbf{q}_k = \mathbf{n}_{k-1}^T \mathbf{q}_{k-1} * (\lambda_{k-1}/\lambda_k) > c_{k-1}$. Thus $\mathbf{q}_k \notin \mathcal{R}_k^0$. \square

Proof of Proposition 1

Proof. By Lemma 1, $\text{diam}(\mathcal{R}_k^0) = 2\sqrt{r_k^2 - (\mathbf{n}_{k-1}^T \mathbf{q}_k - c_{k-1})^2}$. Set $s_k = 1/\lambda_k$. Then,

$$r_k^2 = \|s_k \mathbf{y} - \hat{\boldsymbol{\theta}}_{k-1}\|_2^2 = \|(s_k - s_{k-1})\mathbf{y} + s_{k-1}\mathbf{y} - \hat{\boldsymbol{\theta}}_{k-1}\|_2^2 = \|\Delta_k \mathbf{y} + \beta_{k-1} \mathbf{n}_{k-1}\|_2^2,$$

where $\Delta_k = s_k - s_{k-1}$ and β_{k-1} is a scalar such that $\beta_{k-1} \mathbf{n}_{k-1} = s_{k-1} \mathbf{y} - \hat{\boldsymbol{\theta}}_{k-1}$. Hence

$$r_k^2 = \Delta_k^2 \mathbf{y}^T \mathbf{y} + 2\beta_{k-1} \Delta_k \mathbf{y}^T \mathbf{n}_{k-1} + \beta_{k-1}^2.$$

For the second term we have

$$\begin{aligned} (\mathbf{n}_{k-1}^T \mathbf{q}_k - c_{k-1})^2 &= (\mathbf{n}_{k-1}^T (s_{k-1} \mathbf{y} + s_k \mathbf{y} - s_{k-1} \mathbf{y}) - c_{k-1})^2 \\ &= (\mathbf{n}_{k-1}^T (\Delta_k \mathbf{y} + \beta_{k-1} \mathbf{n}_{k-1}))^2 \\ &= \Delta_k^2 \mathbf{y}^T \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T \mathbf{y} + 2\beta_{k-1} \Delta_k \mathbf{y}^T \mathbf{n}_{k-1} + \beta_{k-1}^2. \end{aligned}$$

Thus $\text{diam}(\mathcal{R}_k^0) = 2\Delta_k \sqrt{\mathbf{y}^T (I - \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T) \mathbf{y}}$. □

Proof of Theorem 1

Proof. Under the DASS feedback rule,

$$s_k = s_1 + 1/2R \sum_{j=2}^k (\mathbf{y}^T (I - \mathbf{n}_{j-1} \mathbf{n}_{j-1}^T) \mathbf{y})^{-1/2} \geq s_1 + (k-1)R/(2\|\mathbf{y}\|_2). \quad (14)$$

Hence $s_1 + (k-1)1/2R/\|\mathbf{y}\|_2 \geq 1/\lambda_t$ is sufficient for termination, from which we deduce $N \leq 1 + (2\|\mathbf{y}\|_2/R)(1/\lambda_t - 1/\lambda_1)$. This bound is very loose. Proposition 1 and the feedback rule (10) give $\text{diam}(\mathcal{R}_k) \leq R$.

The stronger bound on N requires the follow more detailed argument. Let the decreasing sequence $\{\lambda_k\}_{k=1}^N$ be determined by (10) and $\alpha_k = \lambda_k/\lambda_{k-1}$. Then

$$\frac{1}{\lambda_k} - \frac{1}{\lambda_{k-1}} = \left(\frac{1}{\alpha_k} - 1\right) \frac{1}{\lambda_{k-1}} = 1/2R (\mathbf{y}^T (I - \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T) \mathbf{y})^{-1/2}$$

So

$$\alpha_k = \frac{1}{1 + 1/2R \lambda_{k-1} (\mathbf{y}^T (I - \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T) \mathbf{y})^{-1/2}}$$

Using the definition of \mathbf{n}_{k-1} , we have

$$\mathbf{y}^T \mathbf{n}_{k-1} = \frac{\|\mathbf{y}\|_2^2 / \lambda_{k-1} - \mathbf{y}^T \hat{\boldsymbol{\theta}}_{k-1}}{\|\mathbf{y} / \lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1}\|_2} = \frac{\|\mathbf{y}\|_2^2 - \lambda_{k-1} \mathbf{y}^T \hat{\boldsymbol{\theta}}_{k-1}}{\|\mathbf{y} - \lambda_{k-1} \hat{\boldsymbol{\theta}}_{k-1}\|_2}.$$

Hence

$$\mathbf{y}^T (I - \mathbf{n}_{k-1} \mathbf{n}_{k-1}^T) \mathbf{y} = \|\mathbf{y}\|_2^2 - (\mathbf{y}^T \mathbf{n}_{k-1})^2 \quad (15)$$

$$= \|\mathbf{y}\|_2^2 - \frac{(\|\mathbf{y}\|_2^2 - \lambda_{k-1} \mathbf{y}^T \hat{\boldsymbol{\theta}}_{k-1})^2}{\|\mathbf{y} - \lambda_{k-1} \hat{\boldsymbol{\theta}}_{k-1}\|_2^2} \quad (16)$$

$$= \frac{\lambda_{k-1}^2 (\|\mathbf{y}\|_2^2 \|\hat{\boldsymbol{\theta}}_{k-1}\|_2^2 - (\mathbf{y}^T \hat{\boldsymbol{\theta}}_{k-1})^2)}{\|\mathbf{y} - \lambda_{k-1} \hat{\boldsymbol{\theta}}_{k-1}\|_2^2} \quad (17)$$

Let $\mathbf{u}_{k-1} = \hat{\boldsymbol{\theta}}_{k-1} / \|\hat{\boldsymbol{\theta}}_{k-1}\|_2$. We first consider the possibility that $\mathbf{u}_{k-1} = \beta \mathbf{y}$ for some scalar $\beta \neq 0$. So \mathbf{y} and \mathbf{u}_{k-1} are aligned. Since $\lambda_{k-1} < \lambda_{\max}$, in this case we must have $\hat{\boldsymbol{\theta}}_{k-1} = \mathbf{y} / \lambda_{\max}$. Then $\hat{\boldsymbol{\theta}}_k = \mathbf{y} / \lambda_{\max}$ for all k . So this case is easily handled. Hence we will assume below that $\hat{\boldsymbol{\theta}}_{k-1}$ is not aligned with \mathbf{y} .

We now make the assumption that there exists $C > 0$ such that $\|\hat{\boldsymbol{\theta}}(\lambda)\|_2 \leq C$ for $\lambda > 0$. Under this assumption we have:

$$\lambda_{k-1}(\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y})^{-1/2} = \left(\frac{\|\mathbf{y} - \lambda_{k-1}\hat{\boldsymbol{\theta}}_{k-1}\|_2^2}{\|\mathbf{y}\|_2^2\|\hat{\boldsymbol{\theta}}_{k-1}\|_2^2 - (\mathbf{y}^T\hat{\boldsymbol{\theta}}_{k-1})^2} \right)^{1/2} \quad (18)$$

$$\geq \frac{1}{\|\hat{\boldsymbol{\theta}}_{k-1}\|_2} \left(\frac{\|\mathbf{y} - \mathbf{u}_{k-1}\mathbf{u}_{k-1}^T\mathbf{y}\|_2^2}{\|\mathbf{y}\|_2^2 - (\mathbf{y}^T\mathbf{u}_{k-1})^2} \right)^{1/2} \quad (19)$$

$$= \frac{1}{\|\hat{\boldsymbol{\theta}}_{k-1}\|_2} \quad (20)$$

$$\geq \frac{1}{C}. \quad (21)$$

Equation (19) follows from (18) by optimizing over λ_{k-1} . This gives an upper bound on α_k

$$\alpha_k = \frac{1}{1 + 1/2R\lambda_{k-1}(\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y})^{-1/2}} \leq \frac{C}{C + 1/2R} < 1.$$

Set $\alpha = C/(C + 1/2R)$. Then $\lambda_t = \lambda_N = (\prod_{j=2}^N \alpha_j)\lambda_1$. Hence $\log(\lambda_t/\lambda_1) \leq (N - 1)\log \alpha$. Noting that $\lambda_t/\lambda_1 < 1$ and $\alpha < 1$, this implies that

$$N \leq 1 + \frac{\log(\lambda_1/\lambda_t)}{\log(1/\alpha)} \leq 1 + \frac{\log(1/\lambda_t)}{\log(1 + 1/2R/C)}$$

To derive this tighter bound we needed the existence of $\alpha < 1$ with $\alpha_k \leq \alpha$. To show this, we assumed the dual regularization path is bounded. \square

Lemma 3. *If \mathbf{y} is in the range of D , then $\hat{\boldsymbol{\theta}}(\lambda)$ is bounded.*

Proof. By the optimality of $\hat{\mathbf{w}}(\lambda)$, $\mathbf{0}$ must be in the subdifferential of $\|\hat{\mathbf{w}}(\lambda)\|_1$. Hence $D^T(\mathbf{y} - D\hat{\mathbf{w}}(\lambda)) \in \lambda\partial\|\hat{\mathbf{w}}(\lambda)\|_1$. Rearrangement gives $D^T\hat{\boldsymbol{\theta}}(\lambda) \in \partial\|\hat{\mathbf{w}}(\lambda)\|_1$. Let $D = U\Sigma V^T$ be a compact SVD. Then $U^T\hat{\boldsymbol{\theta}}(\lambda) \in \Sigma^{-1}V^T\partial\|\hat{\mathbf{w}}(\lambda)\|_1$. Since \mathbf{y} in the range of D , so is $\hat{\boldsymbol{\theta}}(\lambda)$. Hence $\hat{\boldsymbol{\theta}}(\lambda) \in U\Sigma^{-1}V^T\partial\|\hat{\mathbf{w}}(\lambda)\|_1$. The result follows by observing that each $g \in \partial\|\hat{\mathbf{w}}(\lambda)\|_1$ is bounded. \square

Proof of Corollary 2

Proof. 1) Let $\tilde{\boldsymbol{\theta}}_k = \hat{\boldsymbol{\theta}}_k + \boldsymbol{\epsilon}_k$. The inequality (14) still holds. Hence DASS terminates in a finite number of steps.

2) At step k , we have the solution returned $\tilde{\boldsymbol{\theta}}_{k-1} = \hat{\boldsymbol{\theta}}_{k-1} + \boldsymbol{\epsilon}_{k-1}$, with $\|\hat{\boldsymbol{\theta}}_{k-1} - \tilde{\boldsymbol{\theta}}_{k-1}\|_2 = \|\boldsymbol{\epsilon}_{k-1}\|_2 \leq \rho$ (assumption). As discussed earlier, generically $\hat{\boldsymbol{\theta}}_{k-1}$ is not aligned with \mathbf{y} , and for k large, $\boldsymbol{\theta}_{k-1}$ is not aligned with \mathbf{y} as well. Let $\boldsymbol{\nu}_{k-1} = \hat{\boldsymbol{\theta}}_{k-1}/\|\hat{\boldsymbol{\theta}}_{k-1}\|_2$. Substituting $\hat{\boldsymbol{\theta}}_{k-1}$ with $\tilde{\boldsymbol{\theta}}_{k-1}$ in Equation (18) for the error-free case, we have

$$\lambda_{k-1}(\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y})^{-1/2} = \left(\frac{\|\mathbf{y} - \lambda_{k-1}\tilde{\boldsymbol{\theta}}_{k-1}\|_2^2}{\|\mathbf{y}\|_2^2\|\tilde{\boldsymbol{\theta}}_{k-1}\|_2^2 - (\mathbf{y}^T\tilde{\boldsymbol{\theta}}_{k-1})^2} \right)^{1/2} \quad (22)$$

$$\geq \frac{1}{\|\tilde{\boldsymbol{\theta}}_{k-1}\|_2} \left(\frac{\|\mathbf{y} - \boldsymbol{\nu}_{k-1}\boldsymbol{\nu}_{k-1}^T\mathbf{y}\|_2^2}{\|\mathbf{y}\|_2^2 - (\mathbf{y}^T\boldsymbol{\nu}_{k-1})^2} \right)^{1/2} \quad (23)$$

$$= \frac{1}{\|\tilde{\boldsymbol{\theta}}_{k-1}\|_2} \quad (24)$$

$$= \frac{1}{\|\hat{\boldsymbol{\theta}}_{k-1} + \boldsymbol{\epsilon}_{k-1}\|_2} \quad (25)$$

$$\geq \frac{1}{C + \rho}. \quad (26)$$

Set $\alpha' = (C + \rho)/(C + \rho + 1/2R)$. It follows similarly,

$$N \leq 1 + \frac{\log(\lambda_1/\lambda_t)}{\log(1/\alpha')} \leq 1 + \frac{\log(1/\lambda_t)}{\log(1/\alpha')}.$$

It is easy to see $\alpha' > \alpha$. So for inaccurate solutions, the convergence rate is still of order $O(\log(1/\lambda_t))$, but has a larger upper bound. \square

Effects of the choice of R

We show below another empirical evaluation on the effect of R using a new dataset.

YALEBXF: This uses the frontal face images (size 192×168) of the 38 subjects in the extended Yale B face recognition data set [13, 18]. We randomly select $m = 2,000$ of the 2,414 face images as the dictionary and randomly choose \mathbf{y} from the remaining 414 images.




Figure 8: **Effects of the choice of R on DASS performance.** (Left): Average rejection percentage. (Middle): Average speedup. (Right): The average value of N over the 1200 trials.

The trend in FISTA [4] result is similar to that in FeatureSign [17]. We observed that for the datasets we tested, FeatureSign seems to be faster than FISTA, and FISTA is more sensitive to the reduction of features when used in conjunction with a screening rule and thus has greater speedup.

A Feeckback rule for sequential DPP [29]

We can extend our feedback idea to the sequential DPP rule proposed in [29]. Sequential DPP uses the spherical bound $S_k^{\text{dpp}} = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{k-1}\|_2 \leq 1/\lambda_k - 1/\lambda_{k-1}\}$ to one-shot screen the dictionary at step k . The feedback rule for Sequential DPP is $\lambda_k = (1/2R + \lambda_{k-1}^{-1})^{-1}$. Under this rule, it follows that for all (\mathbf{y}, λ_t) , the sequence terminates after a finite number of steps $N \leq 1 + \frac{2(1/\lambda_t - 1/\lambda_1)}{R} \leq 1 + \frac{2/\lambda_t}{R}$, and $\text{diam}(S_k^{\text{dpp}}) \leq R$, $k = 2:N$. Compared with DASS, feedback controlled sequential DPP has a larger upper bound on the number of iterations required.